



**THE ANALYSIS OF A LINK BETWEEN A REMOTE LOCAL AREA
NETWORK AND ITS SERVER RESOURCES**

THESIS

Theresa D. Beaver, Captain, USAF

AFIT/GCS/ENG/05-02

**DEPARTMENT OF THE AIR FORCE
AIR UNIVERSITY**

AIR FORCE INSTITUTE OF TECHNOLOGY

Wright-Patterson Air Force Base, Ohio

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED

The views expressed in this thesis are those of the author and do not reflect the official policy or position of the United States Air Force, Department of Defense, or the U.S. Government.

AFIT/GCS/ENG/05-02

**THE ANALYSIS OF A LINK BETWEEN A REMOTE LOCAL AREA
NETWORK AND ITS SERVER RESOURCES**

THESIS

Presented to the Faculty

Department of Electrical and Computer Engineering

Graduate School of Engineering and Management

Air Force Institute of Technology

Air University

Air Education and Training Command

In Partial Fulfillment of the Requirements for the

Degree of Master of Science

Theresa D. Beaver, BS

Captain, USAF

June 2005

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED

**THE ANALYSIS OF A LINK BETWEEN A REMOTE LOCAL AREA
NETWORK AND ITS SERVER RESOURCES**

Theresa D. Beaver, BS

Captain, USAF

Approved:

/signed/	9 Jun 05
_____ Robert F. Mills, Ph.D. (Chairman)	_____ Date
/signed/	9 Jun 05
_____ Richard A. Raines, Ph.D. (Member)	_____ Date
/signed/	9 Jun 05
_____ Barry E. Mullins, Ph.D. (Member)	_____ Date

Acknowledgments

I would like to thank my advisor, Dr. Robert Mills, for his assistance with my research, Dr. Rusty Baldwin for answering my OPNET questions, and Tim Lacey for assisting me with problems encountered with my lab computer. I am also grateful for the support I received from my fellow students.

Theresa D. Beaver

Table of Contents

	Page
Acknowledgments.....	iv
Table of Contents.....	v
List of Figures.....	viii
List of Tables.....	ix
Abstract.....	xi
I. Introduction.....	1
Background.....	1
Problem Statement.....	1
Research Goal.....	2
Results Overview.....	2
Summary.....	2
II. Literature Review.....	4
Introduction.....	4
Increased Deployments.....	4
Air Force's Server Consolidation Effort.....	6
Related Network Simulation Research.....	9
Summary.....	16
III. Methodology.....	17
Problem Definition.....	17
Goals.....	17
Approach.....	18
System Boundaries.....	19

System Services	21
Workload	21
Performance Metrics	21
Parameters	22
Factors	24
Evaluation Technique	25
Experimental Design	26
Analysis of Results	27
Summary.....	28
IV. Analysis and Results.....	29
Introduction	29
Results of Simulation Scenarios	29
Experiments with all of the FTP Traffic Routed to the MOB.....	29
Experiments with 1/3 of the FTP Traffic Routed to the Remote Server	36
Experiments with 2/3 of the FTP Traffic Routed to the Remote Server	38
Experiments with All of the FTP Traffic Routed to the Remote Server	40
Summary.....	43
V. Conclusions and Recommendations	44
Research Goals	44
Conclusions of Research	44
Significance of Research	44
Recommendations for Future Research.....	45
Appendix.....	47

Bibliography	52
--------------------	----

List of Figures

	Page
1. The Development Cycle	11
2. Network Overview	19
3. Configurations for 10, 30, 50 and 70 User Networks	20

List of Tables

	Page
1. FTP Traffic Routing Levels	25
2. Percent Link Utilization from the Remote Site to the MOB with all FTP Traffic Routed to the MOB Server	30
3. Computation of Effects for Link Utilization from the Remote Site to the MOB with all FTP Traffic Routed to the MOB Server.....	31
4. Percent Link Utilization from the MOB to the Remote Site with all FTP Traffic Routed to the MOB Server.....	32
5. Computation of Effects for Link Utilization from the MOB to the Remote Site with all FTP Traffic Routed to the MOB Server.....	33
6. TCP Delay in Seconds for Experiments in which 100% of the FTP Traffic was Routed to the MOB Server	34
7. FTP Download Response Time in Seconds for Experiments in which 100% of the FTP Traffic was Routed to the MOB Server	35
8. HTTP Page Response Time in Seconds for Experiments in which 100% of the FTP Traffic was Routed to the MOB Server	36
9. Percent Link Utilization from the Remote Site to the MOB with 1/3 of the FTP Traffic Routed to the Remote Server	37
10. Percent Link Utilization from the MOB to the Remote Site with 1/3 of the FTP Traffic Routed to the Remote Server	38

11. Percent Link Utilization from the Remote Site to the MOB with 2/3 of the FTP Traffic Routed to the Remote Server	39
12. Percent Link Utilization from the MOB to the Remote Site with 2/3 of the FTP Traffic Routed to the Remote Server	40
13. Percent Link Utilization from the Remote Site to the MOB with all FTP Traffic Routed to the Remote Server	41
14. Percent Link Utilization from the MOB to the Remote Site with all FTP Traffic Routed to the Remote Server	42
15. E-mail Application Definitions.....	47
16. FTP Application Definitions.....	48
17. HTTP Application Definitions.....	49
18. Web Browsing (Light) Page Properties	50
19. Web Browsing (Light) Server Selection.....	50
20. Web Browsing (Heavy) Page Properties	50
21. Web Browsing (Heavy) Server Selection.....	51

Abstract

As the Air Force transitions to an expeditionary force, the service's ability to provide computer capabilities at remote locations becomes more and more paramount. One way to provide this support is to create a Local Area Network (LAN) in which the workstations are positioned at the deployed location while the servers are maintained at a Main Operating Base (MOB). This saves the military money, because it eliminates the need to purchase and deploy server equipment as well as eliminating the need to deploy personnel to set-up and maintain the servers. There is, however, a tradeoff. As the number of personnel at the deployed location increases and their computing requirements change, the link between the deployed location and the MOB can become saturated causing degraded performance.

This research looks at how the number of personnel at the deployed location and the types of applications they are using affect the link and the overall system performance. It also examines the effects of adding a server to the deployed location. The results of this study show that the network as configured can support up to 30 users. With the addition of an FTP server at the deployed location, the system can handle 50 users. The system was only able to handle 70 users under the lightest application loads. If the network must support over 50 users, more bandwidth is needed between the deployed location and the MOB.

THE ANALYSIS OF A LINK BETWEEN A REMOTE LOCAL AREA NETWORK AND ITS SERVER RESOURCES

I. Introduction

Background

In the 1990s, Air Force Chief of Staff Gen John P. Jumper led the development of the Expeditionary Aerospace Force concept [COR02]. It was developed in an effort to transform the Air Force into a service capable of efficiently exerting its power anywhere around the globe. The concept is still evolving and is lauded as being the method with which the Air Force is going to achieve its Vision 2020 of “Global Vigilance, Reach, and Power” [AFV00]. In order to support this evolution, the Air Force communications community will have to develop the ability to provide computer support to any point on the globe. The success of the Air Force’s future missions will depend on its ability to provide speedy and reliable computer support to the warfighter regardless of his location.

Problem Statement

Although there are several ways to provide computer support to deployed personnel, this research focuses on a scenario in which the forces are deployed without the servers typically used in a Local Area Network (LAN). The workstations at the deployed location are connected to the servers at a Main Operating Base (MOB) via a dedicated communications link with limited bandwidth. For a small number of personnel with minor computing needs, this link is sufficient, but as the number of deployed

personnel increases, as is apt to happen with a deployment, the link becomes saturated, and the LANs performance degrades.

Research Goal

With the scenario presented here, the Air Force saves money by not deploying servers and the personnel required to maintain them, but at some point, the degradation of the LAN begins to outweigh the advantages gained by not sending the servers to the deployed location. This research seeks to investigate how the number of users and their application usage effects the functioning of the system. At what point does system degradation warrant the deployment of a server or an increase in the link's bandwidth?

Results Overview

This study looked at a LAN with workstations at a deployed location in Baghdad, Iraq and servers at a MOB in Florida. The LAN included a dedicated 64 Kbps link between the two locations. It was found that the link was sufficient for 10 users and 30 users, but it did not perform well under all application levels for 50 users. This problem was alleviated through the use of an FTP server at the deployed location. Unfortunately, the server at the deployed location did not reduce the traffic produced by 70 users enough to make that user level an option. If more that 50 users were added to the network, additional bandwidth would be required.

Summary

The remainder of this thesis is organized as follows. Chapter 2 describes some of the literature used for this research. It covers changes in the Air Force environment,

including its transition into an expeditionary force and its movement towards “one network” [CIS01]. It also describes other research conducted that relates the topic covered by this thesis. Chapter 3 describes the methodology used for creating and conducting the experiments completed for this research. Chapter 4 provides the experiment results and gives an analysis of the data. Chapter 5 concludes this report with conclusions and recommendations for future research.

II. Literature Review

Introduction

This chapter provides background information pertinent to the research conducted for this thesis. The section begins with an explanation of why and how the Air Force is changing how it does business in regards to deployments. This is important, because it shows why it is necessary that the Air Force improve the way it establishes and maintains communications at deployed locations. The next section explains the Air Force's efforts to consolidate its server resources, and the final section describes other network simulation research.

Increased Deployments

The Air Force of the 1980s was extremely large compared to today's force. It had numerous bases abroad with forward bases and an extensive supporting infrastructure in place. Now, the Air Force has a third fewer people and two-thirds fewer overseas bases, yet it conducts four times more deployments and must often take its own infrastructure along [COR02]. Initially, the increase in deployments was treated as a unique event, but in the words of Gen Michael E. Ryan, former Air Force Chief of Staff, "They never seemed to go away." Gen Ryan's statement is confirmed by the fact that during the height of Operation Iraqi Freedom, the Air Force deployed almost 55,000 airmen [POS04].

In the mid-1990s, it became obvious that the Air Force had to change its way of doing business if it hoped to continue meeting its deployment responsibilities with a

reduced force. The officer chosen to lead the effort was Gen John P. Jumper, then commander of 9th Air Force. Gen Jumper is now regarded as the father of the Expeditionary Air Force (EAF) concept which was developed to transition the service into an organization capable of handling these new responsibilities [COR02].

In August 1998, the Air Force announced the move to the EAF concept, and in October 1999, the Air Force began using the plan [COR02]. The concept established 10 Aerospace Expeditionary Forces (AEF), two of which will be deployed or on call at any given time. The EAF established a 15-month rotation cycle. Each AEF is eligible to deploy during three months of that cycle. The rest of the 15-month cycle is reserved for recovery, normal operations and exercises, and preparation for the next deployment. Other components of the EAF construct are the Lead Wings, the Air Expeditionary Wings (AEW), and AEF Prime. Lead Wings are responsible for opening expeditionary bases, while AEWs respond to unexpected developments [AFV00]. AEF Prime includes space, national ISR, long range strike, nuclear, and other assets [POS04].

EAF is the method the Air Force intends to use to accomplish the goals set forth in its newest Vision statement. In May 2000, the Joint Chief's published "Joint Vision 2020", and the Air Force followed suite in June 2000 with the release of "Air Force Vision 2020: Global Vigilance, Reach, and Power". Then Secretary of the Air Force F. Whitten Peters and Gen Ryan were closely involved with the development of the new Vision [COR00]. The statement emphasizes the importance of our military's ability to maintain information superiority. It also states that, "We will streamline what we take with us, reducing our forward support footprint by 50 percent" [AFV00]. The Air Force

seeks to reduce its “forward support footprint” by having the deployed forces use links to space systems to “reach back” to bases in the United States for combat support [COR00]. Part of this initiative will include efforts to reduce the amount of communications equipment deployed.

Air Force’s Server Consolidation Effort

In April 2000, Secretary Peters, his chief information officer, Lawrence Delaney, and General Ryan, visited Silicon Valley to meet with Cisco, Microsoft, Oracle, Sun Microsystems, and Hewlett Packard [MUR02]. The purpose of the visit was to collect information which would aid in formulating a plan for improving the Air Force’s computer network. At the time, the Air Force’s IT landscape was decentralized, fragmented, and expensive to maintain. The lack of consistent technical standards and operating procedures resulted in a network that was impossible to effectively manage [THI04].

As a result of this meeting, Secretary Peters and General Ryan invited Cisco, Sun, and Oracle to participate in the Air Force IT Summit in July 2000. The information exchanged between Air Force leadership and the civilian IT professionals at the summit led to the development of the “One Air Force... One Network” strategy [CIS01].

Server consolidation is only part of this plan for transforming the way the Air Force manages information, but it is a significant part. By consolidating servers, the Air Force hopes to reduce the number of servers, achieve a lower ratio of system administrators to end users, and consolidate services in the fewest places possible without

compromising security or service. According to an article printed in the Air Force Communications Agency's Intercom, consolidation will [SER00]:

- Enable the “massing” of experts in fewer locations to increase mission synergy, depth, and breadth of training, systems support, and help desk coverage.
- Reduce system complexity and unplanned variation in kinds of services, configuration, and system maintenance.
- Help improve the information assurance by consolidating network command and control, reducing the number of Internet entry and exit ports, and increasing ability to watch friendly and suspicious activity more closely.

Despite resistance by those apprehensive about relinquishing control of their IT resources, most Air Force Major Commands (MAJCOMs) have already begun server consolidation efforts [MUR02]. According to the Air Force's Network/Server Consolidation Architecture Systems View, control and maintenance of the new system will be handled by the following four tiers:

- Air Force Network Operations Center (AFNOC) and the Air Force Computer Emergency Response Team (AFCERT) at the Air Force level
- Network Operations and Security Centers (NOSCs) at the MAJCOM level
- Network Control Centers (NCCs) at base level
- Base Geographically Separated Units (GSUs)

These four levels of Air Force Network operations are the foundation of the Air Force's Information Assurance program. They provide commanders with the real-time visibility, management, and control of networks that are crucial to information assurance and to Air Force mission success [SYS02].

In order to support the goal of network/server consolidation, the NCC and NOSC must provide first-class facilities with robust backed-up power (POWER), adequate bandwidth (PIPE), and extremely high system availability (PING), also known as the three Ps. The NCC and NOSC must also maintain highly skilled personnel capable of managing the maintained resources. If it is determined that neither capable manpower nor adequate facilities exist to facilitate server consolidation, the MAJCOM shall take action to hire/train personnel, and upgrade facilities as soon as possible to prepare for optimal consolidation of server assets [OPE02].

In November 2004, the Air Force struck a deal with Microsoft and Dell as part of the "One Air Force... One Network" initiative [GAL04]. The organizations entered into a single enterprise agreement worth as much as \$500 million over six years. This multiyear deal will see Microsoft provide core server software, maintenance and upgrade support, while Dell will supply more than 525,000 Microsoft desktop Windows and Office software licenses [GAL04]. Although Air Force leadership has come under some scrutiny for having made a deal with the software company that some believe is responsible for most of the current insecurities in the Air Force network, Air Force CIO John Gilligan supported the agreement, because it will replace 38 decentralized software contracts and nine support contracts. He estimated that the consolidated contracts will

save the Air Force \$100 million over six years, or over roughly 20 percent over existing contracts [BRE04].

According to an article that appeared in Military Information Technology, the Army is pursuing similar consolidation goals. Chief for the Information Infrastructure Modernization Directorate within the Office of the Army Chief Information Officer/G6, Colonel Mark Barnette stated “We need reach back capability. It’s easier to reach back to a consolidated footprint rather than a fragmented footprint of servers. In order to do reach back in an affordable and effective manner, you need to be able to reach back to a consolidated location.” Hal Stern of Sun Microsystems went on to explain that the military’s current network configuration ties applications to a physical infrastructure making it very hard to move capabilities. He said “If you go to a structure where you centralize the access control and essentially dissociate it from an access device, and you want to talk now about mobilizing people to a different building or a different geography, we are going to go build access devices there, whether they are network computers or PCs or something else, and they are still going to come back to this same infrastructure.” Mr. Stern believes that through server consolidation, the military can separate physical device and geography from that of access control by reducing the number of moving parts within the network [MCC04].

Related Network Simulation Research

Network simulation is a popular method for network performance analysis. While analytical modeling is still used, in many cases it presents an over simplistic view of the network and is unable to simulate the dynamic nature of a network. There are

several network simulation tools available. They include NIST Net, REAL, INSANE, NetSim, Maisie, ns-2, VINT, U-Net, USC TCP-Vegas test-bed, Harvard simulator, COMNETIII, and OPNET [CHA99]. OPNET is the network simulation tool used for this research, because it is the DoD sanctioned software program for communication system simulation. For maximum effectiveness, a simulation environment should be modular, hierarchical, and take advantage of the graphical capabilities of today's workstations. OPNET is an object-oriented simulation environment that meets all these requirements and is the most powerful general-purpose network simulator [CHA99]. It uses discrete event simulations as the means of analyzing system performance and behavior. The key features of OPNET are summarized as follows [CHA99]:

- **Modeling and Simulation Cycle.** OPNET provides powerful tools to assist users in building models, executing simulations, and analyzing output data. See Figure 1 for a graphical representation of the portions of the development cycle OPNET is capable to emulating.
- **Hierarchical Modeling.** OPNET employs a hierarchical structure for creating models. Each level of the hierarchy describes different aspects of the complete model being simulated. OPNET provides four tools called editors to develop a representation of a system being modeled. These are the Network, Node, Process, and Parameter Editors, and they are organized in a hierarchical fashion which supports the concept of model reuse.

- Specialized in communication networks. Detailed library models provide support for existing protocols and allow researchers and developers to either modify these existing models or develop new models of their own.
- Automatic simulation generation. OPNET models can be compiled into executable code. An executable discrete event simulation can be debugged or simply executed, resulting in output data.

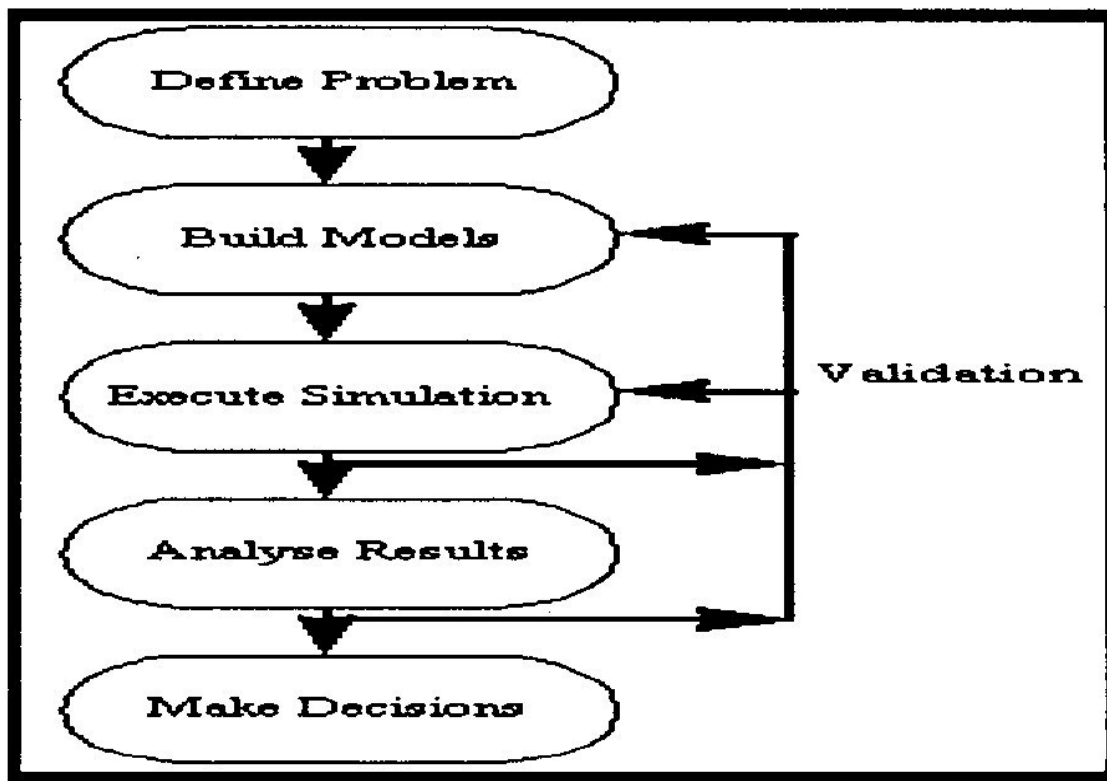


Figure 1. The Development Cycle

OPNET has been used for numerous research projects. For instance, Kishor Waikul used OPNET for several of the projects he conducted while studying at the University of Nevada in Reno. Waikul used OPNET to compare different LAN technologies. With this research, he was able to conclude the following [WAI01]:

- FDDI technology provided minimum delay and gave high throughput because of the minimum bit error rate. Unfortunately, FDDI's high cost may prohibit its use for small enterprises.
- ATM had the highest switching rate and was more suitable for bursty traffic. ATM had comparatively high delay with respect to FDDI, but good stability was ensured.
- Fast Ethernet, which is the most commonly used LAN technology, showed acceptable performance and low cost. It also demonstrated the highest throughput.
- Token ring was slower than the other technologies, but it provided comparable performance at lower cost. It performed best under smaller loads.

Waikul also used OPNET to determine whether or not particular network changes would provide any improvement in system performance. The changes he tested were a change in network topology from star to shared, an upgrade of the links, and migration to a new LAN technology. This study showed that increasing the link capacity increased the speed and hence the amount of data which was circulated within the network. Also, when the size of the network was small in terms of area covered, the shared topology performed better than the star topology. In addition, small networks showed improved performance with the Token Ring topology because the workstations did not have to depend on the central switch for data exchange or link utilization and therefore delay was significantly reduced over the network [WAI02].

In addition to the two studies just mentioned, Waikul also used OPNET to study the effects of adding a firewall to a preexisting network. He simulated the network with the firewall and without the firewall and compared the network performance achieved with the two configurations. The performance characteristics studied were delay, load on the nodes, throughput, and utilization of the links as well as server statistics such as task processing times and page/object response times. Waikul found that introduction of the firewall resulted in increased delay across the network. He speculated that this happened mainly because of the finite IP packet forwarding time taken by the firewall. The throughput was also significantly reduced in some cases. He concluded that, although there was an overall degradation in the network performance, the fact that the data was protected from unauthorized access was worth the cost in system performance [WAI03].

Waikul is of course not the only one conducting research using OPNET. Students at the Center for Satellite and Hybrid Communication Networks at the University of Maryland conducted a study that tested the use of satellites for transmitting TCP/IP traffic [KAR99]. They proposed splitting the TCP connection at the gateways on each end of the satellite link. The model placed the satellite link between segments of the Internet in order to show that the link would not adversely affect the rest of the network. Changes were made to three of the processes within the OPNET router node model to implement the TCP connection splitting. The team studied four simulation scenarios: terrestrial network, satellite network without enhanced gateways, satellite network without enhanced gateway but with large windows used end-to-end, and satellite network with enhanced gateways. The satellite network with the enhanced gateways provided

almost as much throughput at the terrestrial network. In addition, the round trip time measured for the satellite networks without the enhanced gateways was greater than 500ms, while the network with the enhanced gateways had a round trip time of only 75ms facilitating faster recovery from packet loss. File download time was also measured. The enhanced gateways provided file download times 3 to 4 times shorter than the case where the gateways were not used. Although the TCP connection splitting modifications in the gateways did increase the load on the IP layer and reduced the speed at which the gateways could process packets, the team concluded that the network with the enhanced gateways was preferable, because it showed better throughput, shorter measured round trip time, and shorter file download time [KAR99].

Unlike the previous studies which utilized OPNET, another research project of interest was conducted using several other software tools. The purpose of this study was to predict the results of server consolidation. Unlike the system simulated for this thesis, the system under question was already in existence. The study involved an actual bank that acquired a mortgage company and sought to understand the performance impact of centralizing servers and applications. The study predicted there would be degradation in response time for one relocated application, but it also showed that tuning the application would improve the response time without new network investments. In addition, it was shown that moving one of the workloads might overutilize the server, and it could become a performance bottleneck. Further study showed that adding an additional processor to the server alleviated the problem [SPE03].

Despite the fact that this research yielded very interesting results, the primary focus was the method used which the researchers referred to as performance modeling and stepwise refinement. Applying these methods to the server consolidation problem, they produced the following list as a guideline for the research [SPE03]:

- Project planning. This involved verifying consolidation goals and defining model requirements and project milestones.
- Data collection planning. Here, they determined metrics and sources of measurement data.
- Data analysis and profiling. This involved gathering and analyzing the performance characteristics of each transaction as specified by the data collection plan.
- Model creation. This step included developing a simulation model of the application based on the profile.
- Model validation. Here, they ensured model accuracy by comparing model results to actual measured response times.
- Scenario analysis. This included modifying the model to represent the consolidation alternatives, evaluating results, and determining feasibility.

This research, although not as applicable to projects modeling non-existent systems, would be very valuable to anyone looking to consolidate servers for a pre-existing network. It certainly showed the value of determining the effects of a costly system reorganization prior to putting the changes into affect.

Summary

The Air Force is transitioning into an expeditionary force. This transition requires a reevaluation of the means used for supplying computer capabilities to deployed locations. In particular, efforts must be made to reduce the “forward deployed footprint” which includes computer equipment and IT personnel.

At the same time, the Air Force is working to consolidate its servers. Several advantages are gained by collocating and reducing the number of servers supporting the Air Force network. One potential advantage is the creation of a server farm which can be reached by deployed personnel.

A disadvantage is the possibility that placing a link between computers and their servers may cause degradation of the network. One way to reduce this risk is to use network simulations to predict the network performance under particular stressors and to determine bandwidth and server requirements. OPNET has been used for several other research projects to study network performance under various conditions and is the software program chosen for this research.

III. Methodology

Problem Definition

In the early days of computer networking, most nodes of a network were in close vicinity of each other. Over the years, computer networks have become geographically distributed across the world. This paired with the Air Force's transition to an expeditionary force has resulted in the need for a more dynamic and scalable computer network in order to provide reliable communications to forward deployed personnel in a timely manner. The Air Force has also chosen to consolidate its servers. In the not too distant future, each major command will have all its servers collocated at one base, and all the other bases assigned to a major command will have to reach out to that one base for all their server needs. Because of the deployable communication requirements and the decision to consolidate servers, it is now imperative that the Air Force be able to quickly and correctly determine the bandwidth needed to connect users who are geographically separated from their servers. This research investigates how the number of users and the applications being used affect the performance of a connection between a remote location and its server farm.

Goals

The goal of this research is to determine how the number of users at a remote location affects the functioning of the link between the location and its server farm. In addition, the level and type of applications used are examined to determine their impact on the link. And finally, the effects of adding a server to the remote location will be

studied. The hypothesis is that an increase in the number of users or an increase in the use of applications requiring frequent server access will cause a decrease in the effectiveness of the link resulting in longer delays and increased dropped packets. Adding a server to the remote location, should alleviate the stress on the link.

Approach

Two subnets were created in OPNET. One subnet was placed in Baghdad, Iraq, and one was placed near Fort Walton Beach, Florida. The selection of these locations was for illustrative purposes. The results of this research can be applied to any general situation. See Figure 2 for an overview of the two subnets. The subnet in Iraq consists of user workstations in a star configuration attached to a switch via 100baseT Ethernet. This configuration was chosen, because it is the most commonly used configuration. There is also a server attached to this switch via 100baseT Ethernet, but the server is not utilized for all scenarios, because the first set of simulations require all traffic in the network to traverse the transatlantic link. The switch is connected to a router over 100baseT Ethernet. The second subnet consists of a server connected to another router over 100baseT Ethernet. A PPP link with a data rate of 64 Kbps connects the two subnets. Eight different user profiles were created using combinations of high and low email usage, high and low FTP usage, and high and low web browsing. Scenarios were created for 10, 30, 50, and 70 users. Figure 3 shows the configurations for the four different user levels. An extra switch was added to the 70 user configuration, because each switch can only support 64 workstations. During the simulations, link utilization, TCP delay, FTP download response time, and HTTP page response time were recorded.

The values for the individual scenarios will be compared to gain an understanding of how the variables affect the link.

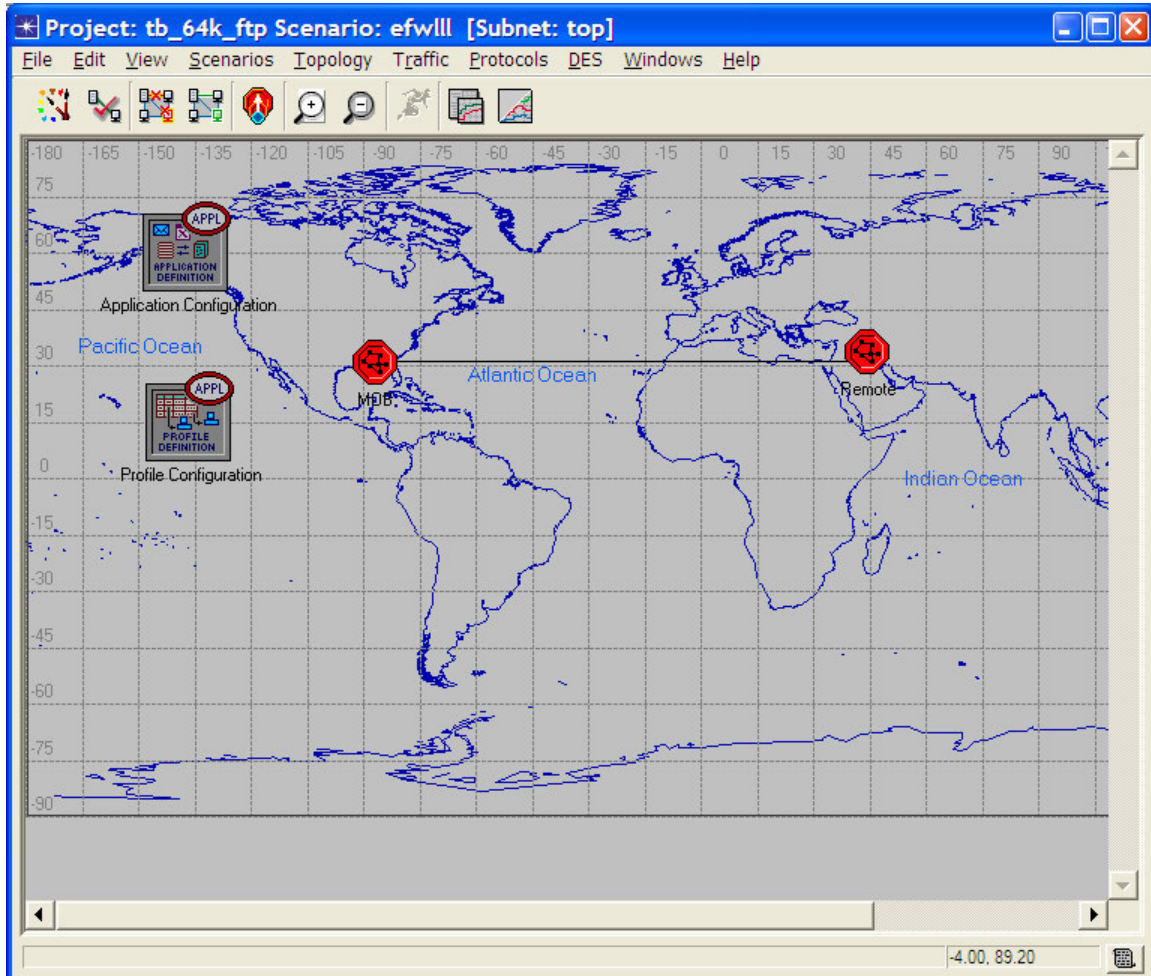


Figure 2. Network Overview

System Boundaries

This section describes the system under test and the component under test. The system under test is a Local Area Network (LAN) in which the workstations are geographically separated from their servers. The system consists of varying numbers of workstations, a switch, two routers, and two servers. The workstations are connected to

the switch in a star topology via 100baseT Ethernet. The switch and the server are attached to their respective routers using 100baseT Ethernet as well. The routers are attached to each other by a Point-to-Point Protocol (PPP) link capable of transmitting 64 Kbps. The key component under study for this research is the link between the router at the workstation location and the router at the server location.




Figure 3. Configurations for 10, 30, 50 and 70 User Networks

System Services

The service this system provides is computer support for personnel geographically separated from their primary computer support providers. The servers at the MOB provide support for e-mail, file transfer, and web browsing.

Workload

The workload introduced to the system was application traffic defined within OPNET. High and low levels of e-mail, file transfer, and web browsing were used. These applications were chosen, because they are the most commonly used. The appendix contains the definitions for the applications used. All eight combinations of the three applications were tested.

Performance Metrics

The performance metrics chosen for this research include link utilization, TCP delay, FTP download response time, and HTTP page response time.

The first performance metric is link utilization. This statistic represents the percentage of the consumption of an available channel bandwidth where a value of 100.0 would indicate full usage.

The second performance metric is TCP delay. The TCP delay is the delay (in seconds) of packets received by the TCP layers in the complete network, for all connections. It is measured from the time an application data packet is sent from the source TCP layer to the time it is completely received by the TCP layer in the destination node.

The third performance metric is FTP download response time. This is the number of seconds elapsed between sending a request and receiving the response packet, and is measured from the time a client application sends a request to the server to the time it receives a response packet. Every response packet sent from a server to an FTP application is included in this statistic.

The final performance metric is HTTP page response time. This specifies the number of seconds required to retrieve an entire page with all the contained inline objects.

Parameters

The parameters for this system include the types of hardware and the link. The hardware was kept simple and constant throughout in order to allow for more realistic comparisons between the different scenarios. Basic Ethernet workstations were used to simulate the users at the remote location. OPNET specifies the Sun Ultra 10 333MHz as the processor for the Ethernet workstation. This was the only processor available for use with the workstations without acquiring additional permission from OPNET.

The links between the workstations and the switch, the switch and the router, and the router and the server are 100BaseT. The 100BaseT duplex link represents an Ethernet connection operation at 100 Mbps. This medium was chosen because of the frequency of its use in actual LANs.

A basic Ethernet switch is used for connecting the workstations at the deployed location. The ethernet64_switch node model represents a switch supporting up to 64 Ethernet interfaces. It implements the Spanning Tree algorithm in order to ensure a loop

free network topology. Packets are received and processed by the switch based on the current configuration of the spanning tree.

The routers used at both the deployed location and the MOB are ethernet4_slip8_gtwy node models. They represent IP-based gateways supporting four Ethernet hub interfaces, and eight serial line interfaces. IP packets arriving on any interface are routed to the appropriate output interface based on their destination IP address. The Routing Information Protocol (RIP) or the Open Shortest Path First (OSPF) protocol may be used to dynamically and automatically create the gateway's routing tables and select routes in an adaptive manner. This gateway requires a fixed amount of time to route each packet as determined by the "IP Routing Speed" attribute of the node. Packets are routed on a first-come-first-serve basis and may encounter queuing at the lower protocol layers, depending on the transmission rate of the corresponding output interface. The routers interfaces include four Ethernet 10BaseT/100BaseT connections and two serial line IP connections at selectable data rates.

The servers at both the MOB and the remote site are simple Ethernet Servers. They represent server nodes with server applications running over TCP/IP and UDP/IP. These nodes support one underlying Ethernet connection at 100 Mbps. The operational speed is determined by the connected link's data rate. A fixed amount of time is required to route each packet as determined by the "IP Forwarding Rate" attribute of the node. Packets are routed on a first-come-first-serve basis and may encounter queuing at the lower protocol layers, depending on the transmission rates of the corresponding output interface. The server is set up to support e-mail, FTP, and web browsing services.

A Point-to-Point Protocol (PPP) link was used to connect the remote site to the MOB. The link is capable of transmitting 64 Kbps and has a propagation speed equal to the speed of light (200,000,000 meters/sec for cable or fiber).

Factors

This research varies the number of users, user application usage, and server location to determine how the link performs under various situations. The number of users was chosen as a factor in order to simulate the increase in personnel that might occur during a deployment. User levels of 10, 30, 50, and 70 were chosen in order to provide a broad, uniform sample.

In order to check the affects of various applications on the link, eight different combinations of application use were tested. High and low levels of e-mail, FTP, and web browsing were used. The combinations are as follows:

- Low e-mail, low FTP, low web browsing (efwlll)
- Low e-mail, low FTP, high web browsing (efwllh)
- Low e-mail, high FTP, low web browsing (efwlhl)
- Low e-mail, high FTP, high web browsing (efwlhh)
- High e-mail, low FTP, low web browsing (efwhll)
- High e-mail, low FTP, high web browsing (efwhlh)
- High e-mail, high FTP, low web browsing (efwhhl)
- High e-mail, high FTP, high web browsing (efwhhh)

The final factor was FTP server location. Servers were set up at both the MOB and the remote site. E-mail and web traffic continued to flow across the transatlantic link. E-mail was not diverted, because it did not have a significant impact on link utilization. Web traffic was not diverted, because such a change in the actual environment represented would require obtaining an Internet Service Provider at the remote location as well as the addition of security measures and equipment at the deployed location. Experiments were run with FTP traffic being directed as shown in Table 1:

Table 1. FTP Traffic Routing Levels

MOB Server	Remote Server
All FTP Traffic	No FTP Traffic
2/3 of FTP Traffic	1/3 of FTP Traffic
1/3 of FTP Traffic	2/3 of FTP Traffic
No FTP Traffic	All FTP Traffic

Evaluation Technique

The three techniques for evaluation are analytical modeling, simulation, and measurement [JAI91]. Simulation was the evaluation technique used in this research. The simulation software package used was OPNET Modeler 10.5. Direct measurement is impossible due to the cost and time required to set up a network similar to the one in the simulation. Analytical modeling wasn't chosen, because the complexity of the problem didn't allow for it.

Experimental Design

The three most frequently used designs are simple designs, full factorial designs, and fractional factorial designs [JAI91]. Full factorial design was used for this research. The advantage of this design is that every possible combination of configuration and workload is examined, and it demonstrates the effect of every factor including the secondary factors and their interactions.

Preliminary experiments were run utilizing varying levels of database, e-mail, and web browsing traffic. The database traffic was later removed and replaced with FTP traffic, because unusual delay values were experienced. After the change, it was discovered that the unusual delay values were a result of flow control mechanisms initiated when the link hit 100% utilization. The FTP traffic was kept, because other research indicated FTP applications are more often used. As discussed in the section on factors, eight workloads were created using high and low levels of e-mail, FTP, and web browsing traffic.

The baseline experiments were run with 10 users at the remote location. All of the FTP traffic was routed to the MOB. All eight combinations of the application traffic were tested. These experiments were run to determine whether or not the simulated network could support 10 users and to show the effects of the applications on link utilization.

Later experiments tested the link with 30, 50, and 70 users. These experiments were run to check the degradation of the link performance as the load increased. This also simulated the effects one might see during the build up phase of a deployment.

The final set of experiments tested the performance of the link as FTP traffic was diverted to a server at the remote location. The results were used to determine at which point it becomes more cost effective to place a server at the remote location. Two of the sets included tests where some of the traffic was sent to the MOB server and some was sent to the remote server. This was done, because it's probable that files would have to be shared between the two locations.

In the end, 128 experiments were run. They break down as follows:

$$8 \text{ workload combinations} \times 4 \text{ user levels} \times 4 \text{ server usage combinations} = \\ 128 \text{ experiments}$$

Each of the 128 experiments was replicated 10 times by using the same 10 seeds for the random number generator. The experiments were replicated 10 times in order to provide more certainty in the averages calculated. The experiments were run to simulate 10 hours of network operation, and 100 values were collected for each variable per experiment yielding a value for every 6 minutes of simulated time. However, the first twelve minutes of simulation time were dropped to eliminate data collected prior to the system reaching steady state. This resulted in 98 data points collected for each experiment.

Analysis of Results

Once all the data was collected, it was used to identify points at which the system performance degraded to such an extent that it became necessary to install a server at the remote location.

Summary

This chapter explained the experiments conducted for this research. The problem was described, and the system and workload were defined. The performance metrics were explained as well as the system parameters and factors. An explanation was given as to why simulation was used for the research, and the full factorial experimental design was presented.

IV. Analysis and Results

Introduction

This chapter focuses on the results and analysis of this research. To reiterate, the purpose of this thesis is to study the effects of adding users to a network in which the workstations are geographically separated from their servers. In addition, changes in the types of applications being used and the effects of rerouting FTP traffic to a server at the deployed locations are studied.

Results of Simulation Scenarios

In all, 128 experiments were performed. In the first 32 experiments, all of the FTP traffic was routed across the transatlantic link to the server at the MOB. These experiments will be used as the baseline to which the other experiments will be compared.

Experiments with all of the FTP Traffic Routed to the MOB

The first metric studied was the link utilization. Because the link is bidirectional, two different sets of data were provided by OPNET, the remote site to the MOB and the MOB to the remote site.

Utilization from the Remote Site to the MOB. Table 2 shows the percentage of link utilization recorded for the indicated experiments. Each of these numbers is an average of 980 different data points recorded by OPNET during 10 repetitions of each experiment. Note that the values collected for 30 users and 50 users are just slightly above 3 times and 5 times higher, respectively, than the values collected for 10 users.

This linear relationship does not carry through for all the application levels for 70 users, because utilization cannot exceed 100%. Notice in the case of high e-mail, high FTP, and high web browsing, 7 times the utilization recorded for 10 users would produce 118.3 percent utilization, which is impossible. Data for scenarios whose traffic exceeds the bandwidth limitations based upon this linear relationship are omitted to reduce confusion. These numbers were of no value, because they were a result of the bandwidth limitation and not a result of the factors being studied.

Table 2. Percent Link Utilization from the Remote Site to the MOB with all FTP Traffic Routed to the MOB Server

	10 Users	30 Users	50 Users	70 Users
Low e-mail, low ftp, low web	0.5	1.7	2.8	3.9
Low e-mail, low ftp, high web	5.5	16.6	27.9	39.5
Low e-mail, high ftp, low web	10.1	29.5	50.1	71.1
Low e-mail, high ftp, high web	14.8	45.0	76.6	-----
High e-mail, low ftp, low web	2.8	8.5	14.2	19.9
High e-mail, low ftp, high web	7.9	23.5	40.5	55.5
High e-mail, high ftp, low web	11.9	36.7	62.0	95.5
High e-mail, high ftp, high web	16.9	52.5	82.3	-----

Table 3 shows the effects of the individual factors for the link utilization for the remote site to the MOB. The data for experiments with 70 users was dropped from these calculations, because the truncation of these values at 100% utilization made them invalid for determining effect. The ‘Row Mean’ values show the average link utilization recorded for the corresponding application combination. The “Column Mean” values

show the average link utilization recorded for the corresponding user levels. The overall average for all scenarios was approximately 26.7 percent. The “Row Effect” values show the difference between the overall average and the average obtained for the corresponding application level. The “Column Effect” values show the difference between the overall average and the average obtained for the corresponding application level. As the table shows, the experiments run using high FTP application usage produced above average link utilization from the remote site to the MOB.

Table 3. Computation of Effects for Link Utilization from the Remote Site to the MOB with all FTP Traffic Routed to the MOB Server

	10 Users	30 Users	50 Users	Row Sum	Row Mean	Row Effect
efwlll	0.55	1.66	2.79	5.00	1.67	-25.03
efwllh	5.53	16.57	27.88	49.97	16.66	-10.04
efwlhl	10.06	29.47	50.11	89.65	29.88	3.18
efwlhh	14.82	45.04	76.59	136.45	45.48	18.78
efwhll	2.82	8.49	14.22	25.54	8.51	-18.19
efwhlh	7.86	23.48	40.50	71.85	23.95	-2.75
efwhhl	11.94	36.67	62.05	110.66	36.89	10.18
efwhhh	16.88	52.49	82.34	151.71	50.57	23.87
Column Sum	70.46	213.86	356.49	640.82		
Column Mean	8.81	26.73	44.56		26.70	
Column Effect	-17.89	0.03	17.86			

Utilization from the MOB to the Remote Site. Table 4 shows the percentage of link utilization recorded for the indicated experiments. As in the previous case, some experiments resulted in link saturation. For instance, for the experiments in which high e-mail, high FTP, and high web browsing were simulated, 5 times the value recorded for the 10 user case produces 136.5 which is not possible. Data for scenarios whose traffic exceeds the bandwidth limitations based upon the linear relationship is omitted to reduce

confusion. The information was of no value, because the values were a result of the bandwidth limitation and not a result of the factors being studied.

A comparison between the link utilization tables shows that there is significantly more traffic being sent from the MOB to the remote site for experiments in which high web browsing is simulated. This occurs because the requests for web pages being sent from the workstations to the server are much smaller than the corresponding web pages being sent from the server back to the workstations.

Table 4. Percent Link Utilization from the MOB to the Remote Site with all FTP Traffic Routed to the MOB Server

	10 Users	30 Users	50 Users	70 Users
Low e-mail, low ftp, low web	0.4	1.2	2.1	2.9
Low e-mail, low ftp, high web	15.8	47.3	78.9	-----
Low e-mail, high ftp, low web	9.6	29.2	48.8	69.2
Low e-mail, high ftp, high web	24.9	75.5	-----	-----
High e-mail, low ftp, low web	2.7	8.1	13.4	18.7
High e-mail, low ftp, high web	18.2	54.1	90.8	-----
High e-mail, high ftp, low web	11.9	35.8	61.4	92.8
High e-mail, high ftp, high web	27.3	82.6	-----	-----

Table 5 shows the effects of the individual factors for the link utilization from the MOB to the remote site. The values for the experiments in which there were 50 and 70 users on the LAN were dropped from these calculations, because the truncation of these values at 100% utilization made them invalid for determining effect. The 'Row Mean' values show the average link utilization recorded for the corresponding

application combination. The “Column Mean” values show the average link utilization recorded for the corresponding user levels. The overall average for all scenarios was approximately 27.8. The “Row Effect” values show the difference between the overall average and the average obtained for the corresponding application level. The “Column Effect” values show the difference between the overall average and the average obtained for the corresponding application level. The experiments run using high web browsing produced above average link utilization from the remote site to the MOB. Notice that this differs from what was seen with the other direction of the link where FTP usage produced above average link utilization.

Table 5. Computation of Effects for Link Utilization from the MOB to the Remote Site with all FTP Traffic Routed to the MOB Server

	10 Users	30 Users	Row Sum	Row Mean	Row Effect
efwlll	0.44	1.25	1.69	0.84	-26.95
efwllh	15.80	47.31	63.11	31.55	3.76
efwlhl	9.61	29.24	38.85	19.43	-8.37
efwlhh	24.92	75.51	100.43	50.22	22.42
efwhll	2.70	8.06	10.76	5.38	-22.41
efwhlh	18.21	54.10	72.31	36.15	8.36
efwhhl	11.87	35.78	47.65	23.82	-3.97
efwhhh	27.29	82.65	109.94	54.97	27.17
Column Sum	110.85	333.89	444.74		
Column Mean	13.86	41.74		27.80	
Column Effect	-13.94	13.94			

TCP Delay. Table 6 shows the TCP delay. The TCP delay is the delay (in seconds) of packets received by the TCP layers in the complete network, for all connections. It is measured from the time an application data packet is sent from the source TCP layer to the time it is completely received by the TCP layer in the destination node. Notice that the system begins to experience high delays in several of the 70-user

experiments and in the 50-user experiments in which high FTP and high web browsing were simulated.

Table 6. TCP Delay in Seconds for Experiments in which 100% of the FTP Traffic was Routed to the MOB Server

	10 Users	30 Users	50 Users	70 Users
Low e-mail, low ftp, low web	0.1	0.1	0.1	0.1
Low e-mail, low ftp, high web	0.3	0.4	0.7	4.9
Low e-mail, high ftp, low web	2.4	3.1	5.1	11.8
Low e-mail, high ftp, high web	0.9	2.5	64.9	84.4
High e-mail, low ftp, low web	0.4	0.4	0.4	0.5
High e-mail, low ftp, high web	0.4	0.5	1.6	13.4
High e-mail, high ftp, low web	1.8	2.7	6.1	99.2
High e-mail, high ftp, high web	0.9	3.5	72.6	81.2

FTP Download Response Time. The FTP download response time is the number of seconds elapsed between sending an FTP request and receiving the response packet. Measured from the time a client application sends a request to the server to the time it receives a response packet. Every response packet sent from a server to an FTP application is included in this statistic. Table 7 shows the FTP download response time for which 100% of the FTP traffic was routed to the MOB server. As the table shows, the FTP download response time becomes unacceptably high when 50 or more users are on the network and are using both FTP and the web browsing applications at high levels. On average, if there were 50 users on the system using high FTP and web applications, it would take approximately 6 ½ minutes to download a 50,000 byte FTP file.

Table 7. FTP Download Response Time in Seconds for Experiments in which 100% of the FTP Traffic was Routed to the MOB Server

	10 Users	30 Users	50 Users	70 Users
Low e-mail, low ftp, low web	0.4	0.5	0.7	0.9
Low e-mail, low ftp, high web	0.6	1.1	2.6	21.7
Low e-mail, high ftp, low web	7.9	11.6	19.7	42.4
Low e-mail, high ftp, high web	9.1	25.4	389.5	484.0
High e-mail, low ftp, low web	0.5	0.8	1.1	1.3
High e-mail, low ftp, high web	0.7	1.5	5.7	52.0
High e-mail, high ftp, low web	8.0	14.5	28.5	411.2
High e-mail, high ftp, high web	9.3	33.4	452.6	496.8

HTTP Page Response Time. The HTTP page response time is the number of seconds required to retrieve an entire page with all the contained inline objects. Table 8 shows the HTTP page response times for experiments in which 100% of the FTP traffic was routed to the MOB server. Notice that the scenarios in which 50 or 70 users are utilizing both FTP and web browsing applications produced extremely high HTTP page response times. As Table 8 shows, 50 users using low e-mail, high FTP, and high web browsing resulted in an HTTP page response time of 145.1 seconds, and 50 users using high e-mail, high FTP, and high web browsing resulted in an HTTP page response time of 163.8 seconds. This means, if there were 50 users on the system using high FTP and web applications, it would take approximately 2 ½ minutes to receive a web page.

Table 8. HTTP Page Response Time in Seconds for Experiments in which 100% of the FTP Traffic was Routed to the MOB Server

	10 Users	30 Users	50 Users	70 Users
Low e-mail, low ftp, low web	0.6	0.7	0.8	1.0
Low e-mail, low ftp, high web	1.4	1.9	3.8	30.2
Low e-mail, high ftp, low web	1.0	2.3	5.1	11.2
Low e-mail, high ftp, high web	2.0	7.7	145.1	172.6
High e-mail, low ftp, low web	0.6	0.8	1.1	1.4
High e-mail, low ftp, high web	1.4	2.2	7.8	72.2
High e-mail, high ftp, low web	1.1	2.8	7.5	103.8
High e-mail, high ftp, high web	2.0	11.0	163.8	175.1

Experiments with 1/3 of the FTP Traffic Routed to the Remote Server

After the scenarios were run in which all the FTP traffic was routed to the MOB server, the next set of scenarios was run with 1/3 of the FTP traffic routed to the server at the remote site. The premise behind this set up was that some of the file storage and exchange could take place at the remote site, but there may still be a need to maintain the ability to use an FTP server at the MOB.

Utilization from the Remote Site to the MOB. Table 9 shows the percentage of link utilization recorded for the indicated experiments in which 1/3 of the FTP traffic was routed to the remote server. The numbers collected for scenarios whose traffic exceeds the bandwidth limitations based upon the linear relationship are omitted to reduce confusion. These numbers were of no value, because they were a result of the bandwidth limitation and not a result of the factors being studied.

Table 9. Percent Link Utilization from the Remote Site to the MOB with 1/3 of the FTP Traffic Routed to the Remote Server

	10 Users	30 Users	50 Users	70 Users
Low e-mail, low ftp, low web	0.5	1.6	2.7	3.8
Low e-mail, low ftp, high web	5.5	16.5	27.7	39.4
Low e-mail, high ftp, low web	7.2	20.8	34.4	47.6
Low e-mail, high ftp, high web	11.7	34.8	58.6	72.7
High e-mail, low ftp, low web	2.8	8.5	14.2	19.8
High e-mail, low ftp, high web	7.8	23.4	40.5	55.2
High e-mail, high ftp, low web	8.8	27.8	42.7	66.9
High e-mail, high ftp, high web	14.3	41.9	75.7	-----

Utilization from the MOB to the Remote Site. Table 10 shows the percentage of link utilization from the MOB to the remote site for the scenarios in which 1/3 of the FTP traffic was routed to the remote server. The information collected for scenarios whose traffic exceeds the bandwidth limitations based upon the linear relationship between the number of users and link utilization are omitted to reduce confusion. The data collected under these circumstances was of no value, because these numbers were a result of the bandwidth limitation and not a result of the factors being studied. Unfortunately, the reduction in traffic obtained by routing 1/3 of the FTP traffic to the remote server was not enough to make putting 50 users on the network a viable option, because the link was still being saturated in cases of high FTP usage and high web browsing.

Table 10. Percent Link Utilization from the MOB to the Remote Site with 1/3 of the FTP Traffic Routed to the Remote Server

	10 Users	30 Users	50 Users	70 Users
Low e-mail, low ftp, low web	0.4	1.2	2.0	2.8
Low e-mail, low ftp, high web	15.8	47.2	78.7	-----
Low e-mail, high ftp, low web	6.9	19.8	33.5	46.3
Low e-mail, high ftp, high web	21.8	65.4	-----	-----
High e-mail, low ftp, low web	2.7	8.0	13.3	18.7
High e-mail, low ftp, high web	18.2	54.1	91.0	-----
High e-mail, high ftp, low web	8.5	27.3	41.5	65.7
High e-mail, high ftp, high web	24.3	72.0	-----	-----

Experiments with 2/3 of the FTP Traffic Routed to the Remote Server

The next set of scenarios tested routed 2/3 of the FTP traffic to the remote server. The premise with these experiments was similar to the last set of experiments. Some of the file storage and exchange could take place at the remote site, but there may still be a need to maintain the ability to use an FTP server at the MOB.

Utilization from the Remote Site to the MOB. The following table shows the percent utilization of the link from the remote site to the MOB for the scenarios in which 2/3 of the FTP traffic was routed to the remote server. Although, all the user and application levels meet the bandwidth requirements the 2/3 scenarios did not bode as well in the other direction on the link.

Table 11. Percent Link Utilization from the Remote Site to the MOB with 2/3 of the FTP Traffic Routed to the Remote Server

	10 Users	30 Users	50 Users	70 Users
Low e-mail, low ftp, low web	0.5	1.6	2.7	3.8
Low e-mail, low ftp, high web	5.5	16.5	27.7	39.3
Low e-mail, high ftp, low web	3.5	11.8	18.9	25.1
Low e-mail, high ftp, high web	8.6	26.8	45.2	59.1
High e-mail, low ftp, low web	2.8	8.4	14.0	19.7
High e-mail, low ftp, high web	7.8	23.3	40.1	55.1
High e-mail, high ftp, low web	6.1	16.4	30.2	42.8
High e-mail, high ftp, high web	10.7	34.1	58.0	73.8

Utilization from the MOB to the Remote Site. Table 12 shows the percent link utilization from the MOB to the remote site for the scenarios in which 2/3 of the FTP traffic was routed to the remote server. Data for scenarios whose traffic exceeds the bandwidth limitations based upon the linear relationship are omitted to reduce confusion. The numbers collected were of no value, because they were a result of the bandwidth limitation and not a result of the factors being studied. Notice that, although the reduction in traffic gained by rerouting 2/3 of the FTP traffic to the remote site server was enough to keep the traffic from the remote site to the MOB at acceptable levels, the same was not true on the link from the MOB to the remote site.

Table 12. Percent Link Utilization from the MOB to the Remote Site with 2/3 of the FTP Traffic Routed to the Remote Server

	10 Users	30 Users	50 Users	70 Users
Low e-mail, low ftp, low web	0.4	1.2	2.0	2.8
Low e-mail, low ftp, high web	15.8	47.3	78.6	-----
Low e-mail, high ftp, low web	3.3	10.9	18.1	23.9
Low e-mail, high ftp, high web	18.9	57.2	93.0	-----
High e-mail, low ftp, low web	2.7	8.0	13.2	18.5
High e-mail, low ftp, high web	18.1	54.0	90.5	-----
High e-mail, high ftp, low web	6.2	16.1	29.7	41.8
High e-mail, high ftp, high web	20.9	65.0	-----	-----

Experiments with All of the FTP Traffic Routed to the Remote Server

The final set of experiments run simulated sending all of the FTP traffic to the remote server. The premise with these experiments was that the personnel at the remote site needed to share files among themselves, but the need to store and share files at the MOB had been eliminated. These experiments showed favorable results for the 50-user configuration.

Utilization from the Remote Site to the MOB. Table 13 shows the percent link utilization from the remote site to the MOB for scenarios in which all of the FTP traffic was routed to the remote server. As with the experiments in which 2/3 of the FTP traffic was routed to the remote server, all user and application levels were within the bandwidth limitations.

Table 13. Percent Link Utilization from the Remote Site to the MOB with all FTP Traffic Routed to the Remote Server

	10 Users	30 Users	50 Users	70 Users
Low e-mail, low ftp, low web	0.5	1.6	2.6	3.6
Low e-mail, low ftp, high web	5.5	16.5	27.6	39.2
Low e-mail, high ftp, low web	0.5	1.6	2.6	3.7
Low e-mail, high ftp, high web	5.5	16.5	27.6	39.2
High e-mail, low ftp, low web	2.8	8.4	14.0	19.5
High e-mail, low ftp, high web	7.9	23.2	40.2	55.1
High e-mail, high ftp, low web	2.8	8.4	14.0	19.5
High e-mail, high ftp, high web	7.7	23.4	40.3	55.2

Utilization from the MOB to the Remote Site. Table 14 shows the percent link utilization from the MOB to the remote site for scenarios in which all of the FTP traffic was routed to the remote server. Data for scenarios whose traffic exceeds the bandwidth limitations based upon the linear relationship between the number of users on the network and the link utilization are omitted to reduce confusion. The numbers collected were of no value, because they were a result of the bandwidth limitation and not a result of the number of users, application usage, or server location. Unlike the previous sets of experiments, now all 50 user application levels meet the bandwidth requirements. The next question is are the other metrics now at acceptable levels.

Table 14. Percent Link Utilization from the MOB to the Remote Site with all FTP Traffic Routed to the Remote Server

	10 Users	30 Users	50 Users	70 Users
Low e-mail, low ftp, low web	0.4	1.1	1.9	2.7
Low e-mail, low ftp, high web	15.7	47.2	78.4	-----
Low e-mail, high ftp, low web	0.4	1.2	1.9	2.7
Low e-mail, high ftp, high web	15.8	47.3	78.6	-----
High e-mail, low ftp, low web	2.6	8.0	13.2	18.4
High e-mail, low ftp, high web	18.1	53.9	90.6	-----
High e-mail, high ftp, low web	2.6	7.9	13.2	18.5
High e-mail, high ftp, high web	18.0	54.0	90.8	-----

TCP Delay. The scenarios which simulated high e-mail, high FTP, and high web browsing produced the highest link utilizations and the highest TCP delays. For this reason, the 50 user configurations using this application level will be compared. In the case where all of the FTP traffic was sent to the MOB server, the TCP delay was 72.6 seconds. Whereas in the case where all of the FTP traffic was routed to the remote server, the TCP delay was only 2.0 seconds. By routing the traffic to the server at the remote site, the TCP delay was reduced by more than one minute.

FTP Download Response Time. In comparing the FTP download response times between the 50 users scenarios using high e-mail, high FTP, and high web browsing, a reduction from 7 ½ minutes to 0.2 seconds is seen. This extreme change is due to the fact that in the later simulations, none of the FTP traffic traverses the transatlantic link.

HTTP Page Response Time. The 50 user high e-mail, high FTP, and high web browsing scenarios yielded an HTTP page response time just under three minutes. Once all the FTP traffic was rerouted to the remote sever, the same configuration provided web pages in 7.8 seconds.

Summary

The evaluation and analysis goals proposed in Chapter 3 were accomplished in this chapter. The effects of the eight application levels and the four user levels were shown as well as the effects of adding an FTP server at the deployed location.

V. Conclusions and Recommendations

Research Goals

The purpose of this research was to characterize the effects of adding users to a network in which the workstations are geographically separated from the servers. In addition, the effects of various application combinations were studied. Later experiments included a server at the remote location which was used to handle varying levels of the networks FTP traffic. The idea was to determine at which point a server would be necessary at the remote location.

Conclusions of Research

OPNET was used to model several scenarios which included four different levels of users and eight different application combinations. The link utilization, TCP delay, FTP download response time, and the HTTP page response time were measured. The experiments showed that the 64 Kbps link was sufficient for all configurations with 10 and 30 users, but the link was never sufficient for configurations with 70 users. The link did not prove sufficient for the 50 user configurations when all the traffic was routed across the link, but by adding an FTP server at the remote location, the traffic crossing the link was reduced just enough to make 50 users a viable option. If more than 50 users were added to the network, a link with more bandwidth would be necessary.

Significance of Research

This research showed the value of OPNET in predicting bandwidth and server requirements. Similar research could be conducted to determine how much bandwidth would be necessary to connect a deployed unit to their in-garrison network or how much

bandwidth would be necessary to connect a unit's network to its parent unit's server farm. In addition, if the traffic patterns of a deploying unit could be determined prior to deployment, more useful studies could be done to predict whether or not it's feasible to deploy the unit without servers, and if so, what changes to the deployed IT environment might make it necessary to deploy servers later during the deployment.

Recommendations for Future Research

The application definitions supplied by OPNET and used for this research are not accurate. Research should be conducted to determine more realistic estimates of the type of traffic produced by the more common computer applications. If such information were collected, it could be used to reconstruct the profiles in the experiments constructed for this research. The results of experiments run using the reconstructed profiles would provide a much more realistic estimate of how well the link in question would function given the various user and application levels.

In addition, The network created for this research should be modified to include a satellite link using the TCP enhanced gateways described in the research conducted at the University of Maryland. Replacing the dedicated Ethernet link of this study with the satellite link of the University of Maryland study will make for a more realistic simulation. In addition, the Maryland researchers suggested their simulation be retested using HTTP, email, and FTP traffic which is the same traffic used for this study [KAR99].

Along another vein, similar research could be conducted to predetermine the results of server consolidation among Air Force units. Such research would provide

valuable information about the bandwidth needed to successfully accomplish consolidation. Researchers interested in attempting such a study should consult the IT Pro paper on server consolidation using performance modeling [SPE03].

Appendix

Table 15. E-mail Application Definitions

	Attribute	Value
Email (Light)	Send Interarrival Time (seconds)	Exponential (3600)
	Send Group Size (seconds)	Constant (3)
	Receive Interarrival Time (seconds)	Exponential (3600)
	Receive Group Size	Constant (3)
	E-Mail Size (bytes)	Constant (500)
	Symbolic Server Name	Email Server
	Type of Service	Best Effort (0)
	RSVP Parameters	None
	Back-End Custom Application	Not Used
Email (Heavy)	Send Interarrival Time (seconds)	Exponential (360)
	Send Group Size (seconds)	Constant (3)
	Receive Interarrival Time (seconds)	Exponential (360)
	Receive Group Size	Constant (3)
	E-Mail Size (bytes)	Constant (2000)
	Symbolic Server Name	Email Server
	Type of Service	Best Effort (0)
	RSVP Parameters	None
	Back-End Custom Application	Not Used

Table 16. FTP Application Definitions

	Attribute	Value
File Transfer (Light)	Command Mix (Get/Total)	50%
	Inter-Request Time (seconds)	Exponential (3600)
	File Size (bytes)	Constant (1000)
	Symbolic Server Name	FTP Server
	Type of Service	Best Effort (0)
	RSVP Parameters	None
	Back-End Custom Application	Not Used
File Transfer (Heavy)	Command Mix (Get/Total)	50%
	Inter-Request Time (seconds)	Exponential (360)
	File Size (bytes)	Constant (50000)
	Symbolic Server Name	FTP Server
	Type of Service	Best Effort (0)
	RSVP Parameters	None
	Back-End Custom Application	Not Used

Table 17. HTTP Application Definitions

	Attribute	Value
Web Browsing (Light)	HTTP Specification	HTTP 1.1
	Page Interarrival Time (seconds)	Exponential (720)
	Page Properties	See Table 16
	Server Selection	See Table 17
	RSVP Parameters	None
	Type of Service	Best Effort (0)
Web Browsing (Heavy)	HTTP Specification	HTTP 1.1
	Page Interarrival Time (seconds)	Exponential (60)
	Page Properties	See Table 18
	Server Selection	See Table 19
	RSVP Parameters	None
	Type of Service	Best Effort (0)

Table 18. Web Browsing (Light) Page Properties

Object Size (bytes)	Number of Objects	Location	Back-End Custom Application	Object Group Name
Constant (500)	Constant (1)	HTTP Server	Not Used	Not Used
Small Image	Constant (5)	HTTP Server	Not Used	Not Used

Table 19. Web Browsing (Light) Server Selection

Attribute	Value
Initial Repeat Probability	Browse
Pages Per Server	Exponential (10)

Table 20. Web Browsing (Heavy) Page Properties

Object Size (bytes)	Number of Objects	Location	Back-End Custom Application	Object Group Name
Constant (1000)	Constant (1)	HTTP Server	Not Used	Not Used
Medium Image	Constant (5)	HTTP Server	Not Used	Not Used

Table 21. Web Browsing (Heavy) Server Selection

Attribute	Value
Initial Repeat Probability	Browse
Pages Per Server	Exponential (10)

Bibliography

- [AFV00] “America’s Air Force Vision 2020”,
<http://www.af.mil/library/posture/vision/vision.pdf>, 2000.
- [BRE04] Brewin, Bob. “Air Force Combines Microsoft Contracts”,
<http://www.fcw.com/fcw/articles/2004/1115/web-afmicro-11-19-04.asp>, November 2004.
- [CHA99] Chang, Xinjie. “Network Simulations with OPNET”,
<http://ieeexplore.ieee.org/iel5/6629/17684/00823089.pdf?isnumber=&arnumber=823089>, 1999.
- [CIS01] “Cisco Systems and the United States Air Force”.
http://www.cisco.com/warp/public/779/ibs/vertical/publicsector/air_force_cp_rev.pdf.
May 2001.
- [COR00] Correll, John T. “Visions”,
http://www.afa.org/magazine/sept2000/0900visions_print.html, September 2000.
- [COR02] Correll, John T. “The EAF in Peace and War”,
<http://www.afa.org/magazine/July2002/0702eaf.html>, July 2002.
- [GAL04] Galli, Peter. “Microsoft, Dell Ink \$500M Deal with Air Force”,
http://www.eweek.com/print_article2/0,2533,a=139646,00.asp, November 2004.
- [JAI91] Jain, Raj. *The Art of Computer Systems Performance Analysis: Techniques for Experimental Design, Measurement, Simulation, and Modeling*. New York: John Wiley & Sons, Inc., 1991.
- [KAR99] Karir, Manish, Mingyan Liu, Bradley Barrett, and John S. Baras. “A Simulation Study of Enhanced TCP/IP Gateways for Broadband Internet over Satellite,” <http://www.isr.umd.edu/CSHCN>, 1999.
- [MCC04] McCarter, Mickey. “Consolidation for Transformation”, <http://www.military-information-technology.com/article.cfm?DocID=715>, December 2004.
- [MUR02] Murray, William. “Air Force Consolidates Network”,
<http://www.esj.com/Features/print.aspx?editorialsId=117>, September 2002.
- [OPE02] AF/XIC. “U.S. Air Force Network/Server Consolidation Architecture, Operational View, High-Level Operational Concept Description”, November 2002.

- [POS04] “United States Air Force Posture Statement 2004”,
<http://www.af.mil/library/posture/posture2004.pdf>, 2004.
- [SER00] “Server Consolidation: Efficiencies from consolidation efforts add up to big savings for Air Force”. <http://public.afca.af.mil/Intercom/2000/nov.pdf>, November 2000.
- [SPE03] Spellman, Amy, Karen Erickson, and Jim Reynolds. “Server Consolidation Using Performance Modeling”,
<http://ieeexplore.ieee.org/iel5/6294/27692/01235607.pdf?isnumber=&arnumber=1235607>, October 2003.
- [SYS02] AF/XIC. “U.S. Air Force Network/Server Consolidation Architecture, Systems View, System Interface Description”, November 2002.
- [THI04] Thibodeau, Patrick. “U.S. Air Force Goes Commercial with Network Consolidation Effort”,
<http://www.computerworld.com/printthis/2004/0,4814,90406,00.html>. March 2005.
- [WAI01] Waikul, Kishor. “Comparison of different LAN Technologies”,
http://www.unr.nevada.edu/~waikul/project_files/Report3_fir.doc.
- [WAI02] Waikul, Kishor. “Enhancing the office network”,
http://www.unr.nevada.edu/~waikul/project_files/Project4_fir.doc.
- [WAI03] Waikul, Kishor. “Study of Firewall Technology & Performance of Network with Firewall”, http://www.unr.nevada.edu/~waikul/project_files/Report_fir.doc.

REPORT DOCUMENTATION PAGE					Form Approved OMB No. 074-0188	
<p>The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of the collection of information, including suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.</p> <p>PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.</p>						
1. REPORT DATE (DD-MM-YYYY)		2. REPORT TYPE		3. DATES COVERED (From – To)		
13-06-2005		Master's Thesis		June 2004 – June 2005		
4. TITLE AND SUBTITLE THE ANALYSIS OF A LINK BETWEEN A REMOTE LOCAL AREA NETWORK AND ITS SERVER RESOURCES				5a. CONTRACT NUMBER		
				5b. GRANT NUMBER		
				5c. PROGRAM ELEMENT NUMBER		
6. AUTHOR(S) Beaver, Theresa D., Captain, USAF				5d. PROJECT NUMBER		
				5e. TASK NUMBER		
				5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAMES(S) AND ADDRESS(S) Air Force Institute of Technology Graduate School of Engineering and Management (AFIT/EN) 2950 Hobson Way, Building 641 WPAFB OH 45433-8865				8. PERFORMING ORGANIZATION REPORT NUMBER AFIT/GCS/ENG/05-02		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) N/A				10. SPONSOR/MONITOR'S ACRONYM(S)		
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)		
12. DISTRIBUTION/AVAILABILITY STATEMENT APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.						
13. SUPPLEMENTARY NOTES						
14. ABSTRACT <p>As the Air Force transitions to an expeditionary force, the service's ability to provide computer capabilities at remote locations becomes more and more paramount. One way to provide this support is to create a Local Area Network (LAN) in which the workstations are positioned at the deployed location while the servers are maintained at a Main Operating Base (MOB). This saves the military money, because it eliminates the need to purchase and deploy server equipment as well as eliminating the need to deploy personnel to set-up and maintain the servers. There is, however, a tradeoff. As the number of personnel at the deployed location increases and their computing requirements change, the link between the deployed location and the MOB can become saturated causing degraded performance.</p> <p>This research looks at how the number of personnel at the deployed location and the types of applications they are using affect the link and the overall system performance. It also examines the effects of adding a server to the deployed location. The results of this study show that the network as configured can support up to 30 users. With the addition of an FTP server at the deployed location, the system can handle 50 users. The system was only able to handle 70 users under the lightest application loads. If the network must support over 50 users, more bandwidth is needed between the deployed location and the MOB.</p>						
15. SUBJECT TERMS Network Analysis, Computer Networks, Simulation, Tactical Communications						
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON	
a. REPORT	b. ABSTRACT	c. THIS PAGE			Robert F. Mills, Ph.D., USAF	
U	U	U	UU	65	19b. TELEPHONE NUMBER (Include area code) (937) 255-6565, ext 4527 (Robert.Mills@afit.edu)	